



US009336094B1

(12) **United States Patent**
Wolfson et al.

(10) **Patent No.:** **US 9,336,094 B1**
(45) **Date of Patent:** **May 10, 2016**

(54) **SCALEOUT REPLICATION OF AN APPLICATION**

(75) Inventors: **Kfir Wolfson**, Beer Sheva (IL); **Assaf Natanzon**, Tel Aviv (IL); **Saar Cohen**, Moshav Mishmeret (IL)

(73) Assignee: **EMC International Company**, Hamilton (BM)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 821 days.

(21) Appl. No.: **13/614,104**

(22) Filed: **Sep. 13, 2012**

(51) **Int. Cl.**
G06F 11/14 (2006.01)
G06F 11/20 (2006.01)
H04L 29/08 (2006.01)
H04L 29/14 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 11/1469** (2013.01); **G06F 11/2025** (2013.01); **G06F 2201/84** (2013.01); **H04L 67/1095** (2013.01); **H04L 69/40** (2013.01)

(58) **Field of Classification Search**
CPC **G06F 2201/84**; **G06F 11/1469**; **G06F 11/2025**; **G06F 11/2089**
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,170,480 A	12/1992	Mohan et al.
5,249,053 A	9/1993	Jain
5,388,254 A	2/1995	Betz et al.
5,499,367 A	3/1996	Bamford et al.
5,526,397 A	6/1996	Lohman
5,864,837 A	1/1999	Maimone
5,879,459 A	3/1999	Gadgil et al.
5,990,899 A	11/1999	Whitten

6,042,652 A	3/2000	Hyun et al.
6,065,018 A	5/2000	Beier et al.
6,143,659 A	11/2000	Leem
6,148,340 A	11/2000	Bittinger et al.
6,174,377 B1	1/2001	Doering et al.
6,174,809 B1	1/2001	Kang et al.
6,203,613 B1	3/2001	Gates et al.
6,260,125 B1	7/2001	McDowell
6,270,572 B1	8/2001	Kim et al.
6,272,534 B1	8/2001	Guha
6,287,965 B1	9/2001	Kang et al.

(Continued)

FOREIGN PATENT DOCUMENTS

EP	1154356	11/2001
WO	WO 00 45581 A3	8/2000

OTHER PUBLICATIONS

Gibson, "Five Point Plan Lies at the Heart of Compression Technology," Apr. 29, 1991; 1 Page.

(Continued)

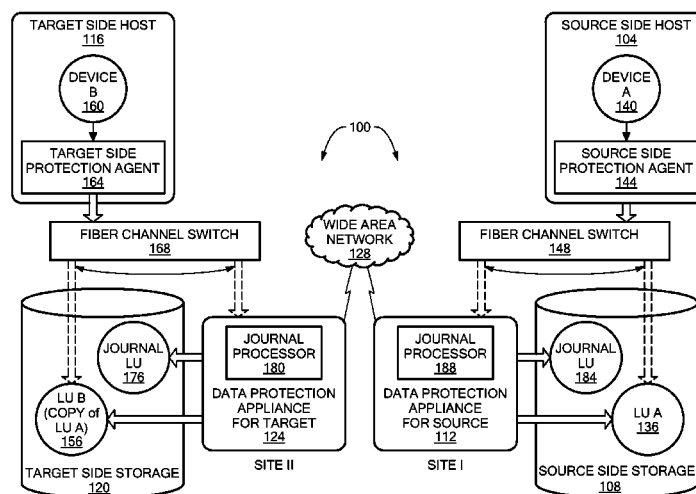
Primary Examiner — Joe Chacko

(74) Attorney, Agent, or Firm — Daly, Crowley, Mofford & Durkee, LLP

(57) **ABSTRACT**

In one aspect, a method includes determining that a first quorum of servers is available at a production site and a target site and generating a group-set bookmark if the first quorum of servers is available. In another aspect, an article includes a non-transitory machine-readable medium that stores executable instructions. The instructions cause a machine to determine that a first quorum of servers is available at a production site and a target site and generate a group-set bookmark if the first quorum of servers is available. In a further aspect, an apparatus includes circuitry configured to determine that a first quorum of servers is available at a production site and a target site; and generate a group -set bookmark if the first quorum of servers available.

13 Claims, 8 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

6,467,023 B1	10/2002	DeKoning et al.	8,601,085 B1	12/2013	Ives et al.
6,574,657 B1	6/2003	Dickinson	8,627,012 B1	1/2014	Derbeko et al.
6,621,493 B1	9/2003	Whitten	8,683,592 B1	3/2014	Dotan et al.
6,804,676 B1	10/2004	Bains, II	8,694,700 B1	4/2014	Natanzon et al.
6,947,981 B2	9/2005	Lubbers et al.	8,706,700 B1	4/2014	Natanzon et al.
7,039,827 B2 *	5/2006	Meyer et al. 714/4.11	8,712,962 B1	4/2014	Natanzon et al.
7,043,610 B2	5/2006	Horn et al.	8,719,497 B1	5/2014	Don et al.
7,051,126 B1	5/2006	Franklin	8,725,691 B1	5/2014	Natanzon
7,076,620 B2	7/2006	Takeda et al.	8,725,692 B1	5/2014	Natanzon et al.
7,111,197 B2	9/2006	Kingsbury et al.	8,726,066 B1	5/2014	Natanzon et al.
7,117,327 B2	10/2006	Hirakawa et al.	8,738,813 B1	5/2014	Natanzon et al.
7,120,768 B2	10/2006	Mizuno et al.	8,745,004 B1	6/2014	Natanzon et al.
7,130,975 B2	10/2006	Suishu et al.	8,751,828 B1	6/2014	Raizen et al.
7,139,927 B2	11/2006	Park et al.	8,769,336 B1	7/2014	Natanzon et al.
7,159,088 B2	1/2007	Hirakawa et al.	8,805,786 B1	8/2014	Natanzon
7,167,963 B2	1/2007	Hirakawa et al.	8,806,161 B1	8/2014	Natanzon
7,203,741 B2	4/2007	Marco et al.	8,825,848 B1	9/2014	Dotan et al.
7,222,136 B1	5/2007	Brown et al.	8,832,399 B1	9/2014	Natanzon et al.
7,296,008 B2	11/2007	Passerini et al.	8,850,143 B1	9/2014	Natanzon
7,328,373 B2	2/2008	Kawamura et al.	8,850,144 B1	9/2014	Natanzon et al.
7,353,335 B2	4/2008	Kawamura	8,862,546 B1	10/2014	Natanzon et al.
7,360,113 B2	4/2008	Anderson et al.	8,892,835 B1	11/2014	Natanzon et al.
7,426,618 B2	9/2008	Vu et al.	8,898,112 B1	11/2014	Natanzon et al.
7,516,287 B2	4/2009	Ahal et al.	8,898,409 B1	11/2014	Natanzon et al.
7,519,625 B2	4/2009	Honami et al.	8,898,515 B1	11/2014	Natanzon
7,519,628 B1	4/2009	Leverett	8,898,519 B1	11/2014	Natanzon et al.
7,546,485 B2	6/2009	Cochran et al.	8,914,595 B1	12/2014	Natanzon
7,577,867 B2	8/2009	Lewin et al.	8,924,668 B1	12/2014	Natanzon
7,590,887 B2	9/2009	Kano	8,930,500 B2	1/2015	Marco et al.
7,606,940 B2	10/2009	Yamagami	8,930,947 B1	1/2015	Derbeko et al.
7,627,612 B2	12/2009	Ahal et al.	8,935,498 B1	1/2015	Natanzon
7,627,687 B2	12/2009	Ahal et al.	8,949,180 B1	2/2015	Natanzon et al.
7,707,184 B1 *	4/2010	Zhang et al. 707/645	8,954,673 B1	2/2015	Natanzon et al.
7,719,443 B1	5/2010	Natanzon	8,954,796 B1	2/2015	Cohen et al.
7,757,057 B2	7/2010	Sangapu et al.	8,959,054 B1	2/2015	Natanzon
7,774,565 B2	8/2010	Lewin et al.	8,977,593 B1	3/2015	Natanzon et al.
7,797,358 B1	9/2010	Ahal et al.	8,977,826 B1	3/2015	Meiri et al.
7,840,536 B1	11/2010	Ahal et al.	8,996,460 B1	3/2015	Frank et al.
7,840,662 B1	11/2010	Natanzon	8,996,461 B1	3/2015	Natanzon et al.
7,844,856 B1	11/2010	Ahal et al.	8,996,827 B1	3/2015	Natanzon
7,849,361 B2	12/2010	Ahal et al.	9,003,138 B1	4/2015	Natanzon et al.
7,860,836 B1	12/2010	Natanzon et al.	9,026,696 B1	5/2015	Natanzon et al.
7,882,286 B1	2/2011	Natanzon et al.	9,031,913 B1	5/2015	Natanzon
7,934,262 B1	4/2011	Natanzon et al.	9,032,160 B1	5/2015	Natanzon et al.
7,958,372 B1	6/2011	Natanzon	9,037,818 B1	5/2015	Natanzon et al.
8,037,162 B2	10/2011	Marco et al.	9,063,994 B1	6/2015	Natanzon et al.
8,041,940 B1	10/2011	Natanzon et al.	9,069,479 B1	6/2015	Natanzon
8,060,713 B1	11/2011	Natanzon	9,069,709 B1	6/2015	Natanzon et al.
8,060,714 B1	11/2011	Natanzon	9,081,754 B1	7/2015	Natanzon et al.
8,103,937 B1	1/2012	Natanzon et al.	9,081,842 B1	7/2015	Natanzon et al.
8,108,634 B1	1/2012	Natanzon et al.	9,087,008 B1	7/2015	Natanzon
8,205,009 B2	6/2012	Hellen et al.	9,087,112 B1	7/2015	Natanzon et al.
8,214,612 B1	7/2012	Natanzon	9,104,529 B1	8/2015	Derbeko et al.
8,250,149 B2	8/2012	Marco et al.	9,110,914 B1	8/2015	Frank et al.
8,271,441 B1	9/2012	Natanzon et al.	9,116,811 B1	8/2015	Derbeko et al.
8,271,447 B1	9/2012	Natanzon et al.	9,128,628 B1	9/2015	Natanzon et al.
8,332,687 B1	12/2012	Natanzon et al.	9,128,855 B1	9/2015	Natanzon et al.
8,335,761 B1	12/2012	Natanzon	9,134,914 B1	9/2015	Derbeko et al.
8,335,771 B1	12/2012	Natanzon et al.	9,135,119 B1	9/2015	Natanzon et al.
8,341,115 B1	12/2012	Natanzon et al.	9,135,120 B1	9/2015	Natanzon
8,370,648 B1	2/2013	Natanzon	9,146,878 B1	9/2015	Cohen et al.
8,380,885 B1	2/2013	Natanzon	9,152,339 B1	10/2015	Cohen et al.
8,392,680 B1	3/2013	Natanzon et al.	9,152,578 B1	10/2015	Saad et al.
8,429,362 B1	4/2013	Natanzon et al.	9,152,814 B1	10/2015	Natanzon
8,433,869 B1	4/2013	Natanzon et al.	9,158,578 B1	10/2015	Derbeko et al.
8,438,135 B1	5/2013	Natanzon et al.	9,158,630 B1	10/2015	Natanzon
8,464,101 B1	6/2013	Natanzon et al.	9,160,526 B1	10/2015	Raizen et al.
8,478,955 B1	7/2013	Natanzon et al.	2002/0129168 A1	9/2002	Kanai et al.
8,495,304 B1	7/2013	Natanzon et al.	2002/0188711 A1 *	12/2002	Meyer et al. 709/223
8,510,279 B1	8/2013	Natanzon et al.	2003/0048842 A1	3/2003	Fourquin et al.
8,521,691 B1	8/2013	Natanzon	2003/0061537 A1	3/2003	Cha et al.
8,521,694 B1	8/2013	Natanzon	2003/0110278 A1	6/2003	Anderson
8,543,609 B1	9/2013	Natanzon	2003/0145317 A1	7/2003	Chamberlain
8,583,885 B1	11/2013	Natanzon	2003/0196147 A1	10/2003	Hirata et al.
8,600,945 B1	12/2013	Natanzon et al.	2004/0205092 A1	10/2004	Longo et al.
			2004/0250032 A1	12/2004	Ji et al.
			2004/0254964 A1	12/2004	Kodama et al.
			2005/0015663 A1 *	1/2005	Armangau et al. 714/15
			2005/0028022 A1	2/2005	Amano

(56)

References Cited

U.S. PATENT DOCUMENTS


2005/0049924	A1	3/2005	DeBettencourt et al.	
2005/0172092	A1	8/2005	Lam et al.	
2005/0273655	A1	12/2005	Chow et al.	
2006/0031647	A1	2/2006	Hirakawa et al.	
2006/0047996	A1	3/2006	Anderson et al.	
2006/0064416	A1	3/2006	Sim-Tang	
2006/0107007	A1	5/2006	Hirakawa et al.	
2006/0117211	A1	6/2006	Matsunami et al.	
2006/0161810	A1	7/2006	Bao	
2006/0179343	A1	8/2006	Kitamura	
2006/0195670	A1	8/2006	Iwamura et al.	
2006/0212462	A1	9/2006	Hellen et al.	
2007/0055833	A1	3/2007	Vu et al.	
2007/0162513	A1	7/2007	Lewin et al.	
2007/0180304	A1	8/2007	Kano	
2007/0198602	A1	8/2007	Ngo et al.	
2007/0198791	A1	8/2007	Iwamura et al.	
2007/0220311	A1*	9/2007	Lewin et al.	714/6
2007/0266053	A1	11/2007	Ahal et al.	
2008/0082591	A1	4/2008	Ahal et al.	
2008/0082592	A1	4/2008	Ahal et al.	
2008/0082770	A1	4/2008	Ahal et al.	

OTHER PUBLICATIONS

Soules, "Metadata Efficiency in Versioning File Systems;" 2003; 16 Pages.

AIX System Management Concepts: Operating Systems and Devices; May 2000; 280 Pages.
 Soules et al.; "Metadata Efficiency in a Comprehensive Versioning File System;" May 2002; CMU-CS-02-145; School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213; 33 Pages.
 Linux Filesystems; Sams Publishing; 2002; 12 Pages.
 Bunyan, "Multiplexing in a BrightStor® ARCserve® Backup Release 11;" Mar. 2004; 4 Pages.
 Marks, "Network Computing;" Feb. 2, 2006; 8 Pages.
 Hill, "Network Computing;" Jun. 8, 2006; 9 Pages.
 Microsoft Computer Dictionary; 2002; Press Fifth Edition; 3 Pages.
 Retrieved from <http://en.wikipedia.org/wiki/DEFLATE>; Deflate; Jun. 19, 2008; 6 Pages.
 Retrieved from http://en.wikipedia.org/wiki/Huffman_coding; Huffman Coding; Jun. 8, 2008; 11 Pages.
 Retrieved from <http://en.wikipedia.org/wiki/LZ77>; LZ77 and LZ78; Jun. 17, 2008; 2 Pages.
 U.S. Appl. No. 11/609,560.
 U.S. Appl. No. 12/057,652.
 U.S. Appl. No. 11/609,561.
 U.S. Appl. No. 11/356,920.
 U.S. Appl. No. 10/512,687.
 U.S. Appl. No. 11/536,233.
 U.S. Appl. No. 11/536,215.
 U.S. Appl. No. 11/536,160.
 U.S. Appl. No. 11/964,168.

* cited by examiner



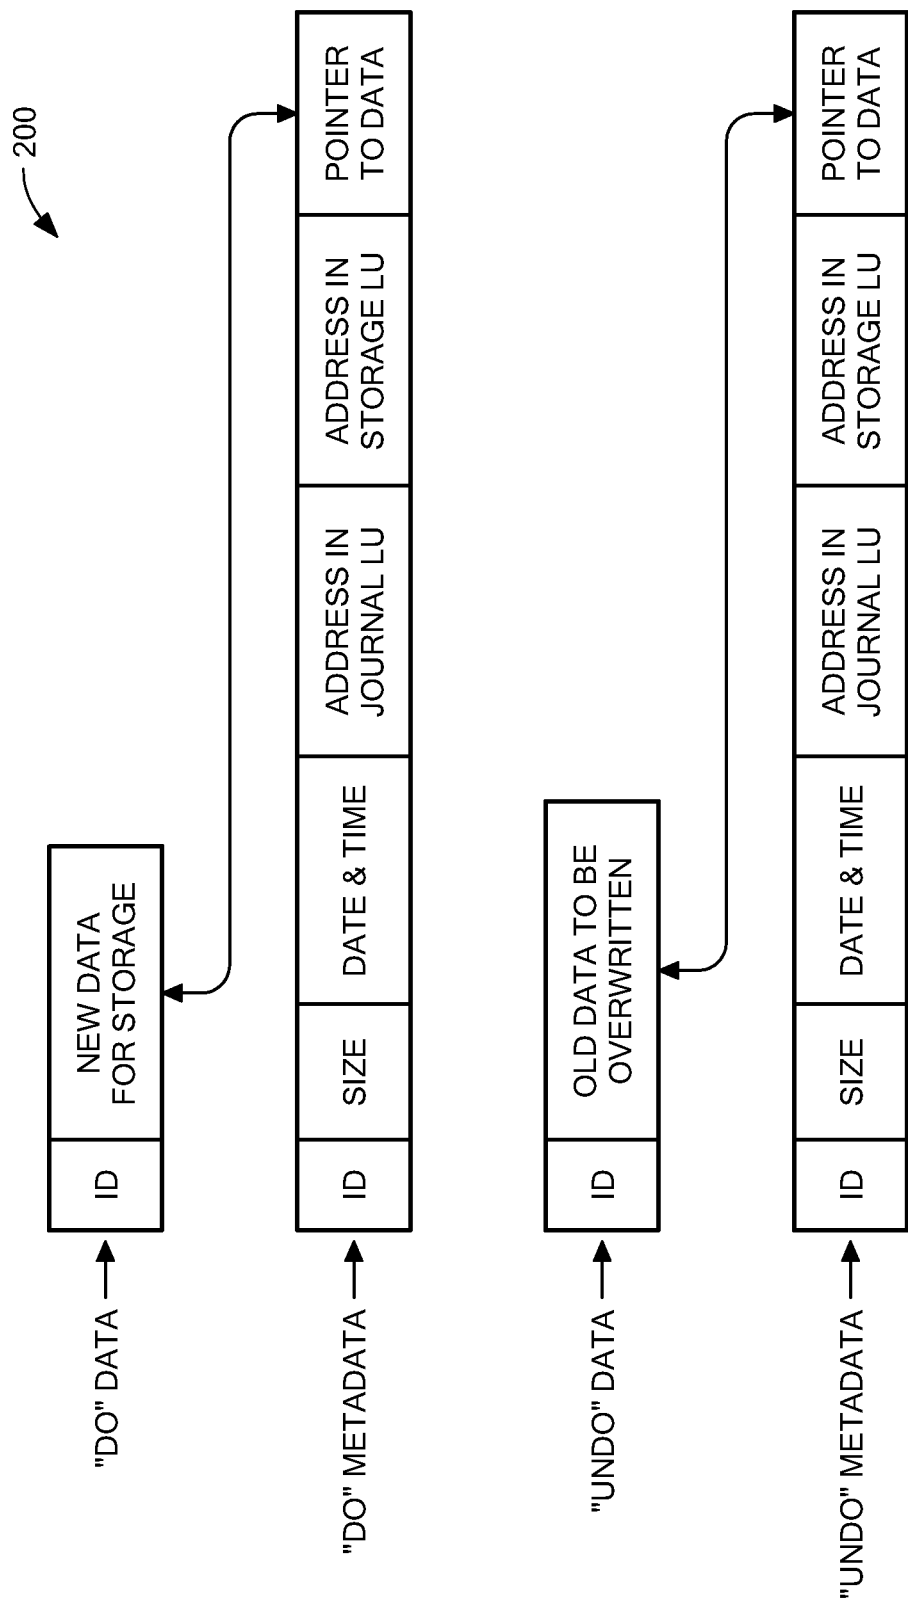


FIG. 2

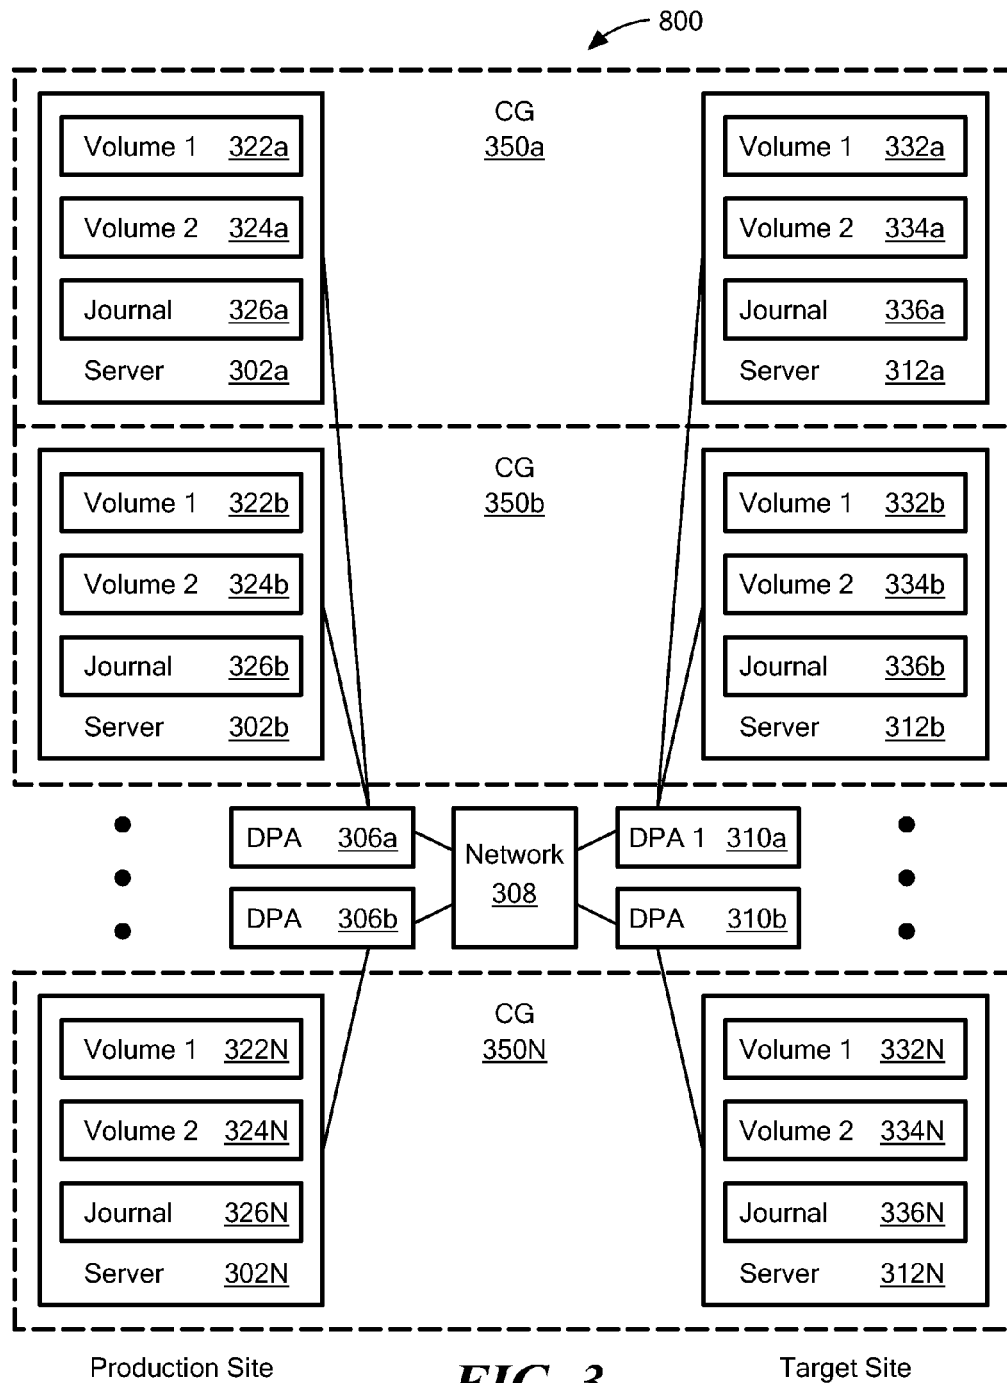
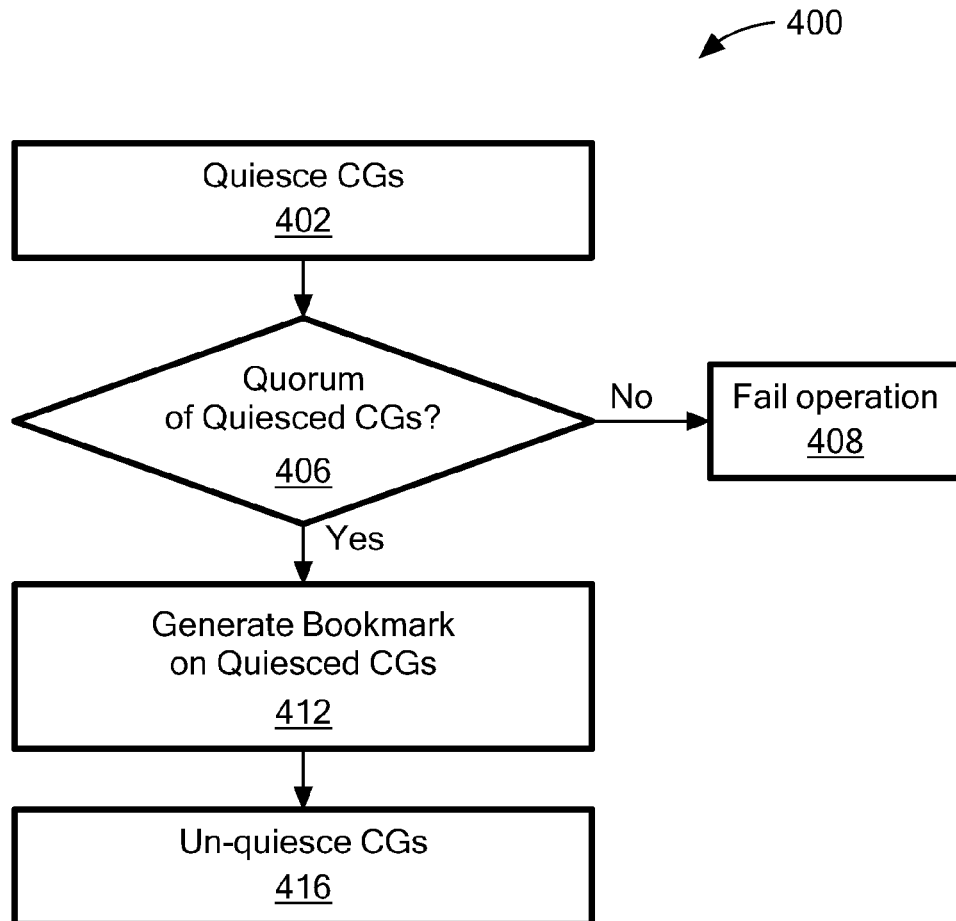
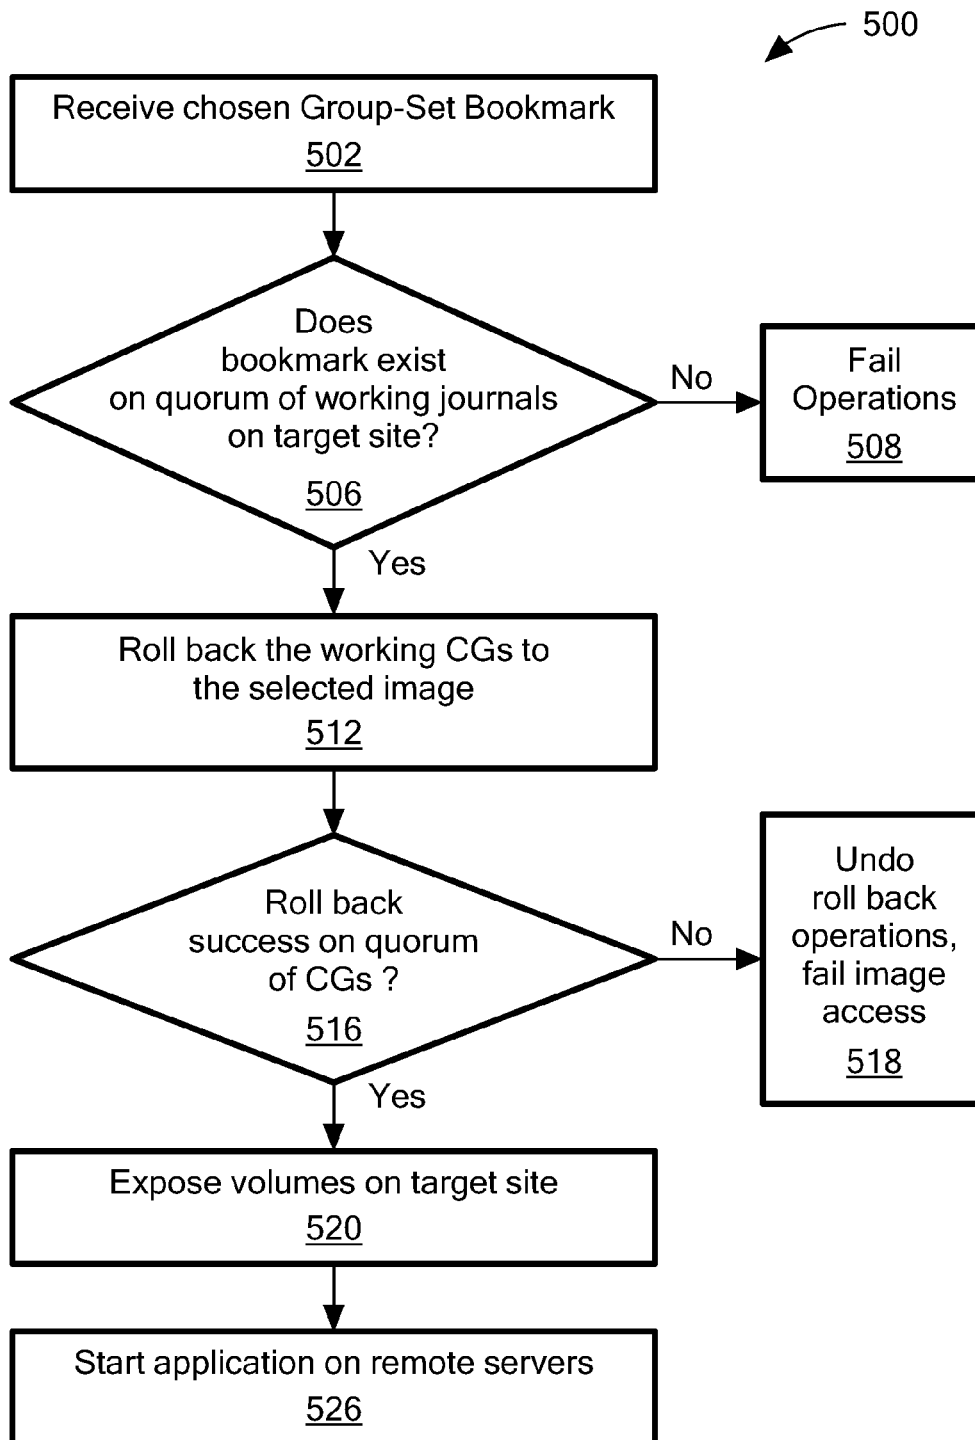
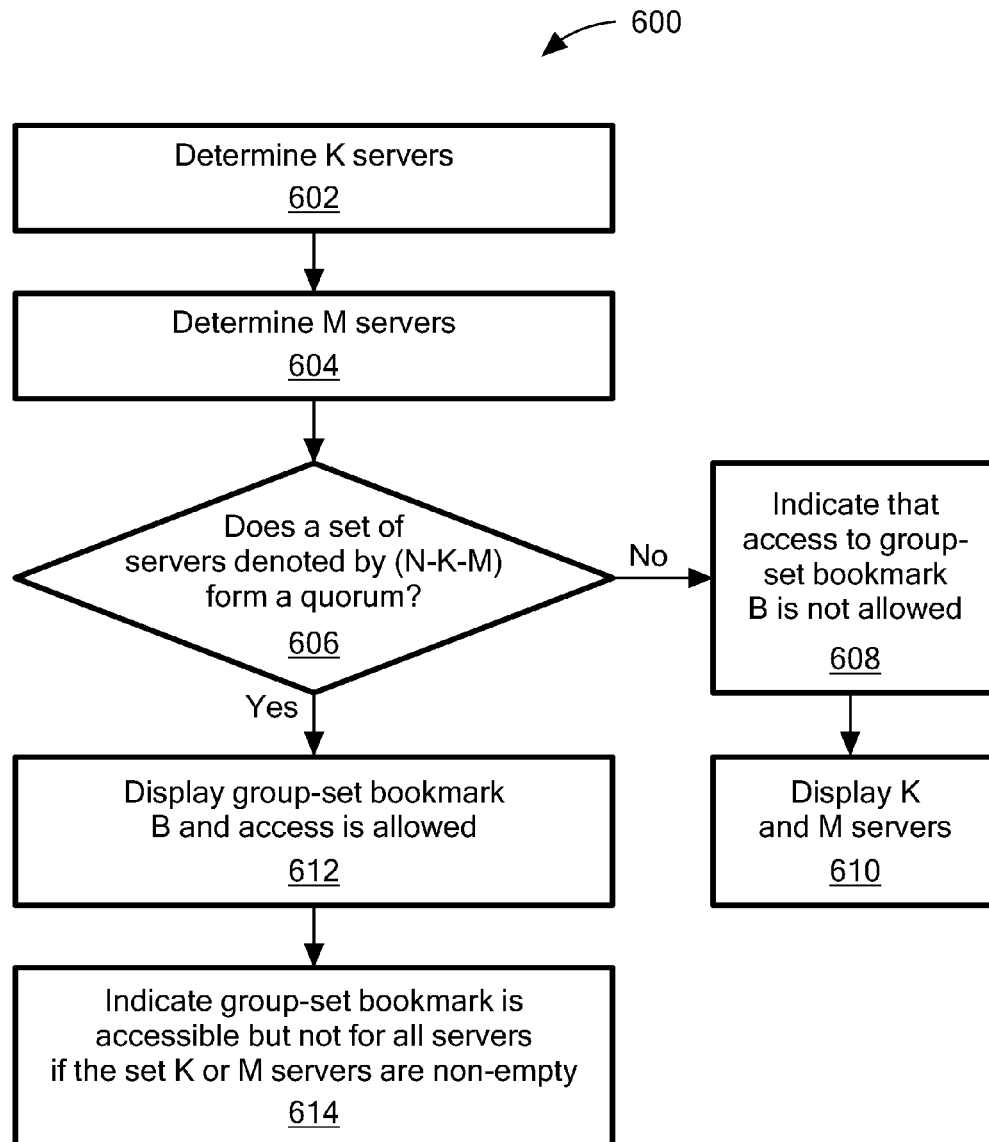
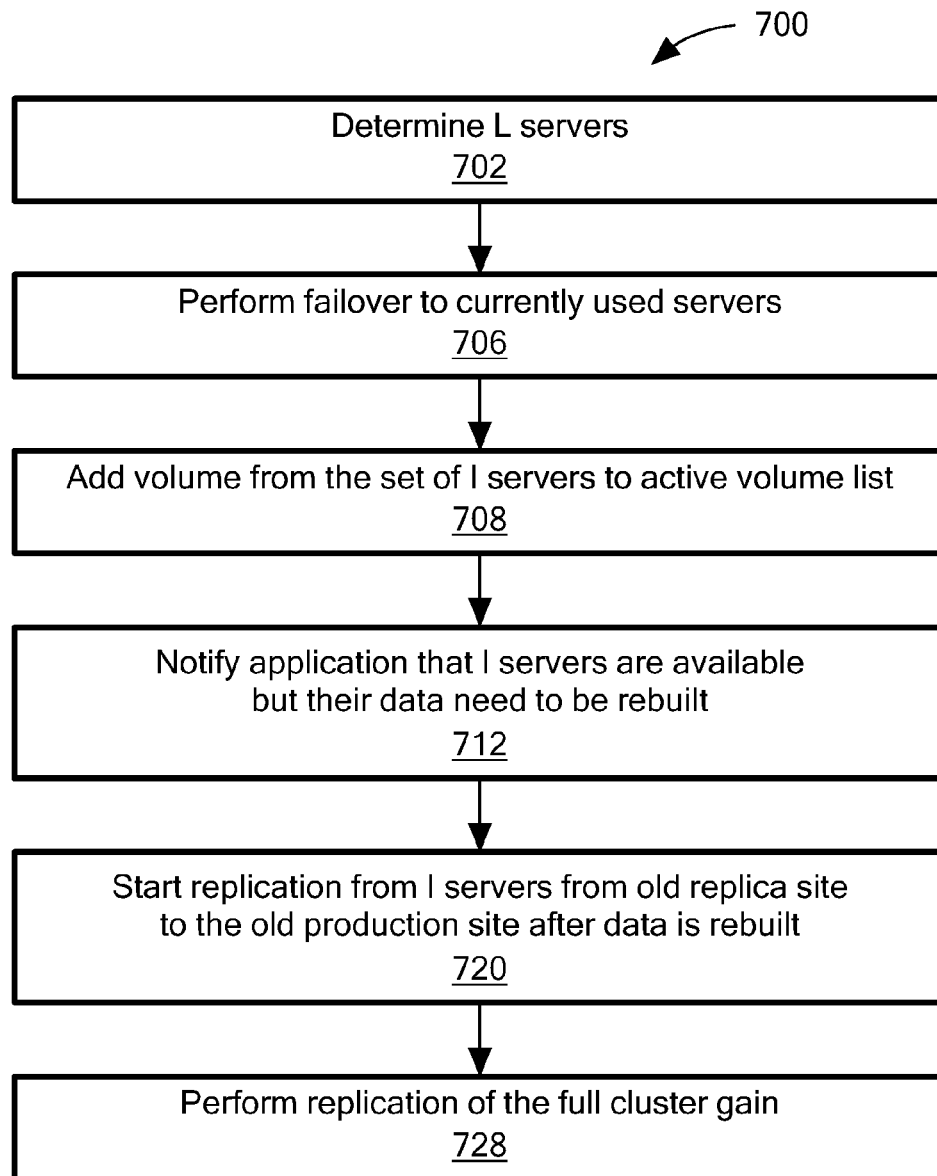


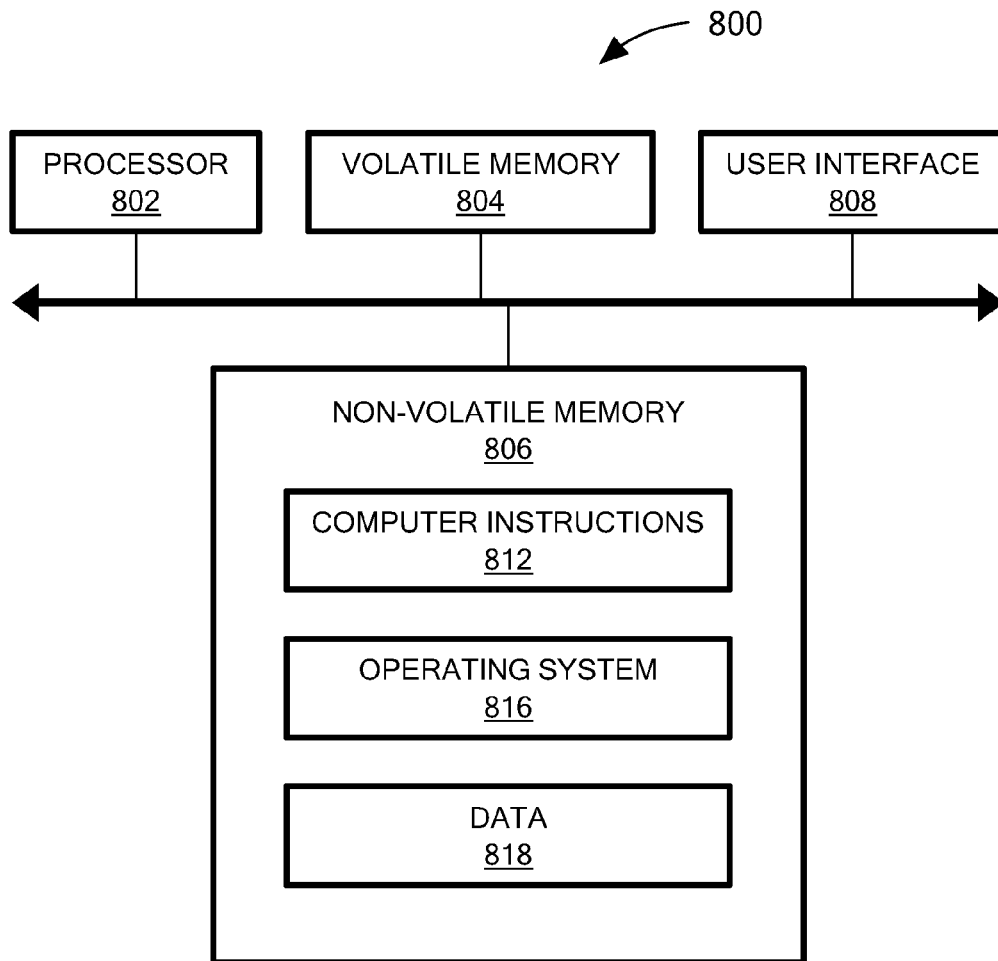
FIG. 3

**FIG. 4**

**FIG. 5**

**FIG. 6**

**FIG. 7**

***FIG. 8***

1

SCALEOUT REPLICATION OF AN APPLICATION

BACKGROUND

Computer data is vital to today's organizations and a significant part of protection against disasters is focused on data protection. As solid-state memory has advanced to the point where cost of memory has become a relatively insignificant factor, organizations can afford to operate with systems that store and process terabytes of data.

Conventional data protection systems include tape backup drives, for storing organizational production site data on a periodic basis. Another conventional data protection system uses data replication, by creating a copy of production site data of an organization on a secondary backup storage system, and updating the backup with changes. The backup storage system may be situated in the same physical location as the production storage system, or in a physically remote location. Data replication systems generally operate either at the application level, at the file system level, or at the data block level.

SUMMARY

In one aspect, a method includes determining that a first quorum of servers is available at a production site and a target site and generating a group-set bookmark if the first quorum of servers is available. In another aspect, an article includes a non-transitory machine-readable medium that stores executable instructions. The instructions cause a machine to determine that a first quorum of servers is available at a production site and a target site and generate a group-set bookmark if the first quorum of servers is available. In a further aspect, an apparatus includes circuitry configured to determine that a first quorum of servers is available at a production site and a target site; and generate a group -set bookmark if the first quorum of servers is available.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of an example of a data protection system.

FIG. 2 is an illustration of an example of a journal history of write transactions for a storage system.

FIG. 3 is a diagram of another example of a data protection system.

FIG. 4 is a flowchart of an example of a process to generate a group-set bookmark.

FIG. 5 is a flowchart of an example of a process to use the group-set bookmark for recovery.

FIG. 6 is a flowchart of an example of a process to list group-set bookmarks.

FIG. 7 is a flowchart of an example of a process to failover.

FIG. 8 is a computer on which any of the processes of FIGS. 4 to 7 may be implemented.

DETAILED DESCRIPTION

Some scale out applications use internal data storage, and provide availability by internal mirroring of the data, for example, data mirroring between segments. When replicating these applications, there is a need to provide scale out replication (i.e., replicate using many replication nodes). Consistency across the replication nodes may be achieved by quiescing the application and generating a bookmark. As described herein, a mechanism is added which allows gen-

2

eration of a bookmark even in cases of a partial disaster in the application nodes. This mechanism allows recovery even if not all of the application nodes are alive at the replica site at the time of the recovery. This is accomplished by introducing knowledge of the availability of the application into data protection agents (splitters) and data protection appliances (DPAs). In the case of a partial disaster at the production site, a bookmark is generated if there is enough available data to recover. At the replica site, a bookmark will allow recovery if there is sufficient information to run the application. Since the application nodes hold the data, disaster of a node means that there is real data loss. However, as described herein, recovery may still occur if there is still enough data to recover.

The following definitions may be useful in understanding the specification and claims.

BACKUP SITE—a facility where replicated production site data is stored; the backup site may be located in a remote site or at the same location as the production site;

BOOKMARK—a bookmark is metadata information stored in a replication journal which indicates a point in time.

DATA PROTECTION APPLIANCE (DPA)—a computer or a cluster of computers responsible for data protection services including inter alia data replication of a storage system, and journaling of I/O requests issued by a host computer to the storage system;

HOST—at least one computer or networks of computers that runs at least one data processing application that issues I/O requests to one or more storage systems; a host is an initiator with a SAN;

HOST DEVICE—an internal interface in a host, to a logical storage unit;

IMAGE—a copy of a logical storage unit at a specific point in time;

INITIATOR—a node in a SAN that issues I/O requests;

I/O REQUEST—an input/output request which may be a read I/O request (read request) or a write I/O request (write request), also referred to as an I/O;

JOURNAL—a record of write transactions issued to a storage system; used to maintain a duplicate storage system, and to roll back the duplicate storage system to a previous point in time;

LOGICAL UNIT—a logical entity provided by a storage system for accessing data from the storage system. The logical disk may be a physical logical unit or a virtual logical unit;

LUN—a logical unit number for identifying a logical unit;

PHYSICAL LOGICAL UNIT—a physical entity, such as a disk or an array of disks, for storing data in storage locations that can be accessed by address;

PRODUCTION SITE—a facility where one or more host computers run data processing applications that write data to a storage system and read data from the storage system;

REMOTE ACKNOWLEDGEMENTS—an acknowledgement from remote DPA to the local DPA that data arrived at the remote DPA (either to the appliance or the journal)

SPLITTER ACKNOWLEDGEMENT—an acknowledgement from a DPA to the protection agent that data has been received at the DPA; this may be achieved by SCSI status cmd.

SAN—a storage area network of nodes that send and receive an I/O and other requests, each node in the network being an initiator or a target, or both an initiator and a target;

SOURCE SIDE—a transmitter of data within a data replication workflow, during normal operation a production site is the source side; and during data recovery a backup site is the source side;

STORAGE SYSTEM—a SAN entity that provides multiple logical units for access by multiple SAN initiators

TARGET—a node in a SAN that replies to I/O requests;

TARGET SIDE—a receiver of data within a data replication workflow; during normal operation a back site is the target side, and during data recovery a production site is the target side;

VIRTUAL LOGICAL UNIT—a virtual storage entity which is treated as a logical unit by virtual machines;

WAN—a wide area network that connects local networks and enables them to communicate with one another, such as the Internet.

A description of journaling and some techniques associated with journaling may be described in the patent titled “METHODS AND APPARATUS FOR OPTIMAL JOURNALING FOR CONTINUOUS DATA REPLICATION” and with U.S. Pat. No. 7,516,287, which is hereby incorporated by reference.

AN EXAMPLE OF A REPLICATION SYSTEM

Referring to FIG. 1, a data protection system **100** includes two sites; Site I, which is a production site and Site II, which is a backup site or replica site. Under normal operation the production site is the source side of system **100**, and the backup site is the target side of the system. The backup site is responsible for replicating production site data. Additionally, the backup site enables roll back of Site I data to an earlier pointing time, which may be used in the event of data corruption of a disaster, alternatively in order to view or to access data from an earlier point in time.

FIG. 1 is an overview of a system for data replication of either physical or virtual logical units. Thus, one of ordinary skill in the art would appreciate that in a virtual environment a hypervisor, in one example, would consume logical units and generate a distributed file system on them such as VMFS creates files in the file system and expose the files as logical units to the virtual machines (each VMDK is seen as a SCSI device by virtual hosts). In another example, the hypervisor consumes a network based file system and exposes files in the NFS as SCSI devices to virtual hosts.

During normal operations, the direction of replicate data flow goes from source side to target side. It is possible, however, for a user to reverse the direction of replicate data flow, in which case Site I starts to behave as a target backup site, and Site II starts to behave as a source production site. Such change of replication direction is referred to as a “failover”. A failover may be performed in the event of a disaster at the production site, or for other reasons. In some data architectures, Site I or Site II behaves as a production site for a portion of stored data, and behaves simultaneously as a backup site for another portion of stored data. In some data architectures, a portion of stored data is replicated to a backup site, and another portion is not.

The production site and the backup site may be remote from one another, or they may both be situated at a common site, local to one another. Local data protection has the advantage of minimizing data lag between target and source, and remote data protection has the advantage of being robust in the event that a disaster occurs at the source side.

The source and target sides communicate via a wide area network (WAN) **128**, although other types of networks may be used.

Each side of system **100** includes three major components coupled via a storage area network (SAN); namely, (i) a storage system, (ii) a host computer, and (iii) a data protection appliance (DPA). Specifically with reference to FIG. 1, the source side SAN includes a source host computer **104**, a source storage system **108**, and a source DPA **112**. Similarly,

the target side SAN includes a target host computer **116**, a target storage system **120**, and a target DPA **124**. As well, the protection agent (splitter) may run on the host, or on the storage, or in the network or at a hypervisor level, and that DPAs are optional and DPA code may run on the storage array too, or the DPA **124** may run as a virtual machine.

Generally, a SAN includes one or more devices referred to as “nodes”. A node in a SAN may be an “initiator” or a “target”, or both. An initiator node is a device that is able to initiate requests to one or more other devices; and a target node is a device that is able to reply to requests, such as SCSI commands, sent by an initiator node. A SAN may also include network switches, such as fiber channel switches. The communication links between each host computer and its corresponding storage system may be any appropriate medium suitable for data transfer, such as fiber communication channel links.

The host communicates with its corresponding storage system using small computer system interface (SCSI) commands.

System **100** includes source storage system **108** and target storage system **120**. Each storage system includes physical storage units for storing data, such as disks or arrays of disks. Typically, storage systems **108** and **120** are target nodes. In order to enable initiators to send requests to storage system **108**, storage system **108** exposes one or more logical units (LU) to which commands are issued. Thus, storage systems **108** and **120** are SAN entities that provide multiple logical units for access by multiple SAN initiators.

Logical units are a logical entity provided by a storage system, for accessing data stored in the storage system. The logical unit may be a physical logical unit or a virtual logical unit. A logical unit is identified by a unique logical unit number (LUN). Storage system **108** exposes a logical unit **136**, designated as LU A, and storage system **120** exposes a logical unit **156**, designated as LU B.

LU B is used for replicating LU A. As such, LU B is generated as a copy of LU A. In one embodiment, LU B is configured so that its size is identical to the size of LU A. Thus for LU A, storage system **120** serves as a backup for source side storage system **108**. Alternatively, as mentioned hereinabove, some logical units of storage system **120** may be used to back up logical units of storage system **108**, and other logical units of storage system **120** may be used for other purposes. Moreover, there is symmetric replication whereby some logical units of storage system **108** are used for replicating logical units of storage system **120**, and other logical units of storage system **120** are used for replicating other logical units of storage system **108**.

System **100** includes a source side host computer **104** and a target side host computer **116**. A host computer may be one computer, or a plurality of computers, or a network of distributed computers, each computer may include inter alia a conventional CPU, volatile and non-volatile memory, a data bus, an I/O interface, a display interface and a network interface. Generally a host computer runs at least one data processing application, such as a database application and an e-mail server.

Generally, an operating system of a host computer creates a host device for each logical unit exposed by a storage system in the host computer SAN. A host device is a logical entity in a host computer, through which a host computer may access a logical unit. Host device **104** identifies LU A and generates a corresponding host device **140**, designated as Device A, through which it can access LU A. Similarly, host computer **116** identifies LU B and generates a corresponding device **160**, designated as Device B.

In the course of continuous operation, host computer **104** is a SAN initiator that issues I/O requests (write/read operations) through host device **140** to LU A using, for example, SCSI commands. Such requests are generally transmitted to LU A with an address that includes a specific device identifier, an offset within the device, and a data size. Offsets are generally aligned to 512 byte blocks. The average size of a write operation issued by host computer **104** may be, for example, 10 kilobytes (KB); i.e., 20 blocks. For an I/O rate of 50 megabytes (MB) per second, this corresponds to approximately 5,000 write transactions per second.

System **100** includes two data protection appliances, a source side DPA **112** and a target side DPA **124**. A DPA performs various data protection services, such as data replication of a storage system, and journaling of I/O requests issued by a host computer to source side storage system data. As explained in detail herein, when acting as a target side DPA, a DPA may also enable roll back of data to an earlier point in time, and processing of rolled back data at the target site. Each DPA **112** and **124** is a computer that includes inter alia one or more conventional CPUs and internal memory.

For additional safety precaution, each DPA is a cluster of such computers. Use of a cluster ensures that if a DPA computer is down, then the DPA functionality switches over to another computer. The DPA computers within a DPA cluster communicate with one another using at least one communication link suitable for data transfer via fiber channel or IP based protocols, or such other transfer protocol. One computer from the DPA cluster serves as the DPA leader. The DPA cluster leader coordinates between the computers in the cluster, and may also perform other tasks that require coordination between the computers, such as load balancing.

In the architecture illustrated in FIG. 1, DPA **112** and DPA **124** are standalone devices integrated within a SAN. Alternatively, each of DPA **112** and DPA **124** may be integrated into storage system **108** and storage system **120**, respectively, or integrated into host computer **104** and host computer **116**, respectively. Both DPAs communicate with their respective host computers through communication lines such as fiber channels using, for example, SCSI commands or any other protocol.

DPAs **112** and **124** are configured to act as initiators in the SAN; i.e., they can issue I/O requests using, for example, SCSI commands, to access logical units on their respective storage systems. DPA **112** and DPA **124** are also configured with the necessary functionality to act as targets; i.e., to reply to I/O requests, such as SCSI commands, issued by other initiators in the SAN, including inter alia their respective host computers **104** and **116**. Being target nodes, DPA **112** and DPA **124** may dynamically expose or remove one or more logical units.

As described hereinabove, Site I and Site II may each behave simultaneously as a production site and a backup site for different logical units. As such, DPA **112** and DPA **124** may each behave as a source DPA for some logical units, and as a target DPA for other logical units, at the same time.

Host computer **104** and host computer **116** include protection agents **144** and **164**, respectively. Protection agents **144** and **164** intercept SCSI commands issued by their respective host computers, via host devices to logical units that are accessible to the host computers. A data protection agent may act on an intercepted SCSI commands issued to a logical unit, in one of the following ways: send the SCSI commands to its intended logical unit; redirect the SCSI command to another logical unit; split the SCSI command by sending it first to the respective DPA; after the DPA returns an acknowledgement, send the SCSI command to its intended logical unit; fail a

SCSI command by returning an error return code; and delay a SCSI command by not returning an acknowledgement to the respective host computer.

A protection agent may handle different SCSI commands, differently, according to the type of the command. For example, a SCSI command inquiring about the size of a certain logical unit may be sent directly to that logical unit, while a SCSI write command may be split and sent first to a DPA associated with the agent. A protection agent may also change its behavior for handling SCSI commands, for example as a result of an instruction received from the DPA.

Specifically, the behavior of a protection agent for a certain host device generally corresponds to the behavior of its associated DPA with respect to the logical unit of the host device. When a DPA behaves as a source side DPA for a certain logical unit, then during normal course of operation, the associated protection agent splits I/O requests issued by a host computer to the host device corresponding to that logical unit. Similarly, when a DPA behaves as a target device for a certain logical unit, then during normal course of operation, the associated protection agent fails I/O requests issued by host computer to the host device corresponding to that logical unit.

Communication between protection agents and their respective DPAs may use any protocol suitable for data transfer within a SAN, such as fiber channel, or SCSI over fiber channel. The communication may be direct, or via a logical unit exposed by the DPA. Protection agents communicate with their respective DPAs by sending SCSI commands over fiber channel.

Protection agents **144** and **164** are drivers located in their respective host computers **104** and **116**. Alternatively, a protection agent may also be located in a fiber channel switch, or in any other device situated in a data path between a host computer and a storage system or on the storage system itself. In a virtualized environment, the protection agent may run at the hypervisor layer or in a virtual machine providing a virtualization layer.

What follows is a detailed description of system behavior under normal production mode, and under recovery mode.

In production mode DPA **112** acts as a source side DPA for LU A. Thus, protection agent **144** is configured to act as a source side protection agent; i.e., as a splitter for host device A. Specifically, protection agent **144** replicates SCSI I/O write requests. A replicated SCSI I/O write request is sent to DPA **112**. After receiving an acknowledgement from DPA **124**, protection agent **144** then sends the SCSI I/O write request to LU A. After receiving a second acknowledgement from storage system **108** host computer **104** acknowledges that an I/O command complete.

When DPA **112** receives a replicated SCSI write request from data protection agent **144**, DPA **112** transmits certain I/O information characterizing the write request, packaged as a "write transaction", over WAN **128** to DPA **124** on the target side, for journaling and for incorporation within target storage system **120**.

DPA **112** may send its write transactions to DPA **124** using a variety of modes of transmission, including inter alia (i) a synchronous mode, (ii) an asynchronous mode, and (iii) a snapshot mode. In synchronous mode, DPA **112** sends each write transaction to DPA **124**, receives back an acknowledgement from DPA **124**, and in turns sends an acknowledgement back to protection agent **144**. Protection agent **144** waits until receipt of such acknowledgement before sending the SCSI write request to LU A.

In asynchronous mode, DPA **112** sends an acknowledgement to protection agent **144** upon receipt of each I/O request, before receiving an acknowledgement back from DPA **124**.

In snapshot mode, DPA 112 receives several I/O requests and combines them into an aggregate “snapshot” of all write activity performed in the multiple I/O requests, and sends the snapshot to DPA 124, for journaling and for incorporation in target storage system 120. In snapshot mode DPA 112 also sends an acknowledgement to protection agent 144 upon request of each I/O request, before receiving an acknowledgement back from DPA 124.

For the sake of clarity, the ensuing discussion assumes that information is transmitted at write-by-write granularity.

While in production mode, DPA 124 receives replicated data of LU A from DPA 112, and performs journaling and writing to storage system 120. When applying write operations to storage system 120, DPA 124 acts as an initiator, and sends SCSI commands to LU B.

During a recovery mode, DPA 124 undoes the write transactions in the journal, so as to restore storage system 120 to the state it was at, at an earlier time.

As described hereinabove, LU B is used as a backup of LU A. As such, during normal production mode, while data written to LU A by host computer 104 is replicated from LU A to LU B, host computer 116 should not be sending I/O requests to LU B. To prevent such I/O requests from being sent, protection agent 164 acts as a target site protection agent for host Device B and fails I/O requests sent from host computer 116 to LU B through host Device B.

Target storage system 120 exposes a logical unit 176, referred to as a “journal LU”, for maintaining a history of write transactions made to LU B, referred to as a “journal”. Alternatively, journal LU 176 may be striped over several logical units, or may reside within all of or a portion of another logical unit. DPA 124 includes a journal processor 180 for managing the journal.

Journal processor 180 functions generally to manage the journal entries of LU B. Specifically journal processor 180 enters write transactions received by DPA 124 from DPA 112 into the journal, by writing them into the journal LU, reads the undo information for the transaction from LU B, updates the journal entries in the journal LU with undo information, applies the journal transactions to LU B, and removes already-applied transactions from the journal.

Referring to FIG. 2, which is an illustration of a write transaction 200 for a journal. The journal may be used to provide an adaptor for access to storage 120 at the state it was in at any specified point in time. Since the journal contains the “undo” information necessary to roll back storage system 120, data that was stored in specific memory locations at the specified point in time may be obtained by undoing write transactions that occurred subsequent to such point in time.

Write transaction 200 generally includes the following fields: one or more identifiers; a time stamp, which is the date & time at which the transaction was received by source side DPA 112; a write size, which is the size of the data block; a location in journal LU 176 where the data is entered; a location in LU B where the data is to be written; and the data itself.

Write transaction 200 is transmitted from source side DPA 112 to target side DPA 124. As shown in FIG. 2, DPA 124 records the write transaction 200 in the journal that includes four streams. A first stream, referred to as a DO stream, includes new data for writing in LU B. A second stream, referred to as an DO METADATA stream, includes metadata for the write transaction, such as an identifier, a date & time, a write size, a beginning address in LU B for writing the new data in, and a pointer to the offset in the DO stream where the corresponding data is located. Similarly, a third stream, referred to as an UNDO stream, includes old data that was

overwritten in LU B; and a fourth stream, referred to as an UNDO METADATA, include an identifier, a date & time, a write size, a beginning address in LU B where data was to be overwritten, and a pointer to the offset in the UNDO stream where the corresponding old data is located.

In practice each of the four streams holds a plurality of write transaction data. As write transactions are received dynamically by target DPA 124, they are recorded at the end of the DO stream and the end of the DO METADATA stream, prior to committing the transaction. During transaction application, when the various write transactions are applied to LU B, prior to writing the new DO data into addresses within the storage system, the older data currently located in such addresses is recorded into the UNDO stream. In some examples, the metadata stream (e.g., UNDO METADATA stream or the DO METADATA stream) and the data stream (e.g., UNDO stream or DO stream) may be kept in a single stream each (i.e., one UNDO data and UNDO METADATA stream and one DO data and DO METADATA stream) by interleaving the metadata into the data stream.

Referring to FIG. 3, a data protection system 300 includes N servers (e.g., a server 302a, a server 302b, . . . , and a server 302N) on a production site coupled to data protection appliances (e.g., a DPA 306a and a DPA 306b). The data protection system 300 also includes N servers (e.g., a server 312a, a server 302b, . . . , and a server 302N) on a target site (replica site) coupled to data protection appliances (e.g., a DPA 310a and a DPA 310b). The DPAs 306a, 306b are coupled to the DPAs 310a, 310b through a network 308 (e.g., a WAN network, a Fibre Channel network and so forth). One server on the production site and one server on the target site form a consistency group. For example, the server 302a and the server 312a are put of a consistency group (CS) 350a, the server 302b, and the server 312b are part of a consistency group (CG) 350b, . . . , the server 302N and the server 312N are part of a consistency group (CG) 350N.

The servers 302a-302N on the production site include a volume 1, a volume 2 and a journal. In particular, the server 302a includes a volume 1 322a, a volume 2 324a and a journal 326a, the server 302b includes a volume 1 322b, a volume 2 324b and a journal 326b, . . . , and the server 302N includes a volume 1 322N, a volume 2 324N and a journal 326N.

The servers 312a-312N on the target site (replica site) include a volume 1, a volume 2 and a journal. In particular, the server 312a includes a volume 1 332a (a replica of volume 1 322a), a volume 2 334a (a replica of volume 2 324a) and a journal 336a, the server 312b includes a volume 1 332b (a replica of volume 1 322b), a volume 2 334b (a replica of volume 2 324b) and a journal 336b, . . . , and the server 312N includes a volume 1 332N (a replica of volume 1 322N), a volume 2 334N (a replica of volume 2 324N) and a journal 336N.

The servers 302a-302N, 312a-312N on each site (production and target site) run a distributed application (e.g., GREENPLUM® Database). By internally mirroring the data between the servers, the application has a built-in redundancy which allows for a subset of the servers to be down, while the application continues to run (e.g., in a lower -performance mode) on a minimum number of servers called a quorum.

Replicating distributed applications at the volume level requires separately replicating each production server’s internal data volumes to its symmetric equivalent on the target site (e.g., the user volumes and journal volumes reside in internal disks and constitute a CG per server pair). Also, replicating distributed applications requires taking bookmarks across all the CGs at once. A Group-Set feature is used herein for this thereby tying all CGs into one consistent logical unit.

Referring to FIG. 4, an example of a process to generate group-set bookmarks is a process 400. Process 400 quiesces consistency groups with timeout t (402) (i.e., the DPAs are configured to delay acknowledging new I/Os arriving to all of the consistency groups which are quiesced). The delay is performed by holding the status to the SCSI commands sent by the data protection agent (splitter) to the DPAs for all CGs.

Process 400 determines if there is a quorum of consistency groups (406), for example, within a predetermined amount of time. In one particular example, if all data protection agents (splitters) answer immediately then there is no need to wait the predetermined amount of time. The quorum is determined by configuring the DPAs to have knowledge of the internal redundancy of the application. In some embodiments there is a quorum, if all hosts except maybe one have acknowledged. In other embodiments redundancy may be higher and even failure of two or more hosts may allow for a quorum. In other embodiments failure of some servers to acknowledge may still allow a quorum while failure of another set of servers may fail the quorum.

If there is not a quorum of consistency groups after a predetermined amount of time, process 400 fails the operation (408).

If there is a quorum of consistency groups, process 400 generates a bookmark on quiesced consistency groups (412). Process 400 un-quiesces the quiesced consistency groups (416).

Referring to FIG. 5, an example of a process to use the group-set bookmark for recovery is a process 500. Process 500 receives a chosen group-set bookmark (502). For example, a group-set bookmark is chosen by a user which was generated using the process 400.

Process 500 determines if the group-set bookmark chosen exists on a quorum of working journals on the target site (506). For example, this quorum is the same as described through configuration which is specific to the application redundancy. If the group-set bookmark chosen does not exist on a quorum of working journals on the target site, process 500 fails the operation (508).

If the group-set bookmark chosen exists on a quorum of working journals on the target site, process 500 roll back the working CGs to the selected image (512).

Process 500 determines if the roll back is successful in time t on a quorum of CGs (516). If the roll back is not successful in time t on a quorum of CGs, process 500 undoes the roll back operations and fails image access (518).

If the roll back is successful in time t on a quorum of CGs, process 500 exposes volumes on the target site (520) and starts the application on target side servers (526).

Referring to FIG. 6, an example of a process to list group-set bookmarks is a process 600. On NCGs, for each group-set bookmark, B process 600 performs the following processing blocks. Process 600 determines K servers (602). K servers are equal to a subset of the servers on the target side with an unavailable journal.

Process 600 determines M servers (604). M servers are equal to a subset of servers on the target site with an accessible journal but the accessible journal does not include the group-set bookmark B.

Process 600 determines if a set of servers which include N servers less the K servers and less the M servers (denoted as "N-K-M servers") form a quorum (606). That is, if the servers which are available contain all the application's data. For example, if the application is redundant to survive failure of two servers, then if $K=2$, $M=0$ or $K \leq 1$, $N \leq 1$ or $M=2$ and $K=0$, the remaining servers will form a quorum and the system can access the group-set bookmark. A DPA is configured

to include logic that has information on which subsets of the servers forms a quorum. The simplest configuration is redundancy (i.e. how many server failures can the cluster survive). If the set of N-K-M servers does not form a quorum, process 600 indicates (e.g., to a user) that the image access to the group-set bookmark, B, is not allowed (608). Process 600 displays the servers K and M (e.g., to a user) (610).

If the set of N-K-M servers does form a quorum, the group-set bookmark is listed in the list of group-set bookmarks for which access is allowed (612). If the sets K or M are non-empty, an indication is shown (614) that the group-set bookmark is accessible but not all servers will be available.

Referring to FIG. 7, an example of a process to failover is a process 700. The process 700 is performed in an image access mode and the user wants to failover the image. Process 700 determines L servers (702). L servers are a set of servers on a target side which are not participating in a current quorum. For example, L is a union of K and M servers, and a set J of servers which are not available at the old production site.

Process 700 performs failover on currently used servers, e.g., replication starts from all volumes within the quorum (N-L) from the replica site and the production site. (706).

Process 700 adds volumes from a set of I servers, which is a subset of the L servers, which are currently available at the target site but not currently part of the quorum, to the active volume list at the target (replica) site (708). Process 700 notifies the application that the set of I servers are available again but data in the internal application of their volumes needs to be rebuilt (712).

Once the application indicates that the data have been rebuilt, process 700 starts replication from servers in the subset of I, which is also available at the old production site, from the did replica site to the old production site (720).

Process 700 performs replication of all available application nodes again (728). For example, if the subset of I servers included one host for which the group-set bookmark was not generated, out of twenty hosts, replication will start from the working nineteen hosts at the target (replica) site to the old production site. The volumes on the host for which the group-set bookmark failed will be exposed at a different point in time, and the application will be notified that the data on that host is dirty. The application will rebuild the data, on that host and then replication will start from the host for which no group-set bookmark was available to its pair host on the production site.

Referring to FIG. 8, a computer 800 includes a processor 802, a volatile memory 804, a non-volatile memory 806 (e.g., hard disk) and a user interface (UI) 808 (e.g., a mouse, a keyboard, a display, touch screen and so forth). The non-volatile memory 806 stores computer instructions 812, an operating system 816 and data 818. In one example, the computer instructions 812 are executed by the processor 802 out of volatile memory 804 to perform all or part of the processes described herein (e.g., processes 400, 500, 600, 700).

The processes described herein (e.g., processes 400, 500, 600, 700) are not limited to use with the hardware and software of FIG. 8; they may find applicability in any computing or processing environment and with any type of machine or set of machines that is capable of running a computer program. The processes described herein may be implemented in hardware, software, or a combination of the two. The processes described herein may be implemented in computer programs executed on programmable computers/machines that each includes a processor, a storage medium or other article of manufacture that is readable by the processor (including volatile and non-volatile memory and/or storage ele-

11

ments), at least one input device, and one or more output devices. Program code may be applied to data entered using an input device to perform any of the processes described herein and to generate output information.

The system may be implemented, at least in part, via a computer program product, (e.g., in a machine-readable storage device), for execution by, or to control the operation of, data processing apparatus (e.g., a programmable processor, a computer, or multiple computers)). Each such program may be implemented in a high level procedural or object-oriented programming language to communicate with a computer system. However, the programs may be implemented in assembly or machine language. The language may be a compiled or an interpreted language and it may be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program may be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a communication network. A computer program may be stored on a storage medium or device (e.g., CD-ROM, hard disk, or magnetic diskette) that is readable by a general or special purpose programmable computer for configuring and operating the computer when the storage medium or device is read by the computer to perform the processes described herein. The processes described herein may also be implemented as a machine-readable storage medium, configured with a computer program, where upon execution, instructions in the computer program cause the computer to operate in accordance with the processes. A non-transitory machine-readable medium may include but is not limited to a hard drive, compact disc, flash memory, non-volatile memory, volatile memory, magnetic diskette and so forth but does not include a transitory signal per se.

The processes described herein are not limited to the specific examples described. For example, the processes **400**, **500**, **600**, **700** are not limited to the specific processing order of FIGS. **4** to **7**, respectively. Rather, any of the processing blocks of FIGS. **4** to **7** may be re-ordered, combined or removed, performed in parallel or in serial, as necessary, to achieve the results set forth above.

The processing blocks (for example, in the processes **400**, **500**, **600**, **700**) associated with implementing the system may be performed by one or more programmable processors executing one or more computer programs to perform the functions of the system. All or part of the system may be implemented as, special purpose logic circuitry (e.g., an FPGA (field-programmable gate array) and/or an ASIC (application-specific integrated circuit)).

Elements of different embodiments described herein may be combined to form other embodiments not specifically set forth above. Other embodiments not specifically described herein are also within the scope of the following claims.

What is claimed is:

1. A method, comprising:

determining that a first quorum of servers is available at a production site and a target site;

generating a group-set bookmark if the first quorum of servers is available;

failing the operation of generating the group-set bookmark if there is not a first quorum of servers available;

determining L servers, L servers indicating a union of a set of servers on the target site with an unavailable journal or a journal that does not include the group-set bookmark and a union of servers at a production site which are not available;

12

determining if a set of N servers less the L servers forms a quorum, N being the set of servers on the target site;

performing failover to currently used servers;

allowing access to available volumes at the target site which are not part of the quorum;

notifying an application that volumes are to be rebuilt;

starting the replication of the volumes which are available also at the production site, wherein starting the replication of the volumes which are available also at the production site comprises replicating from working servers at the target site to the production site;

exposing at a different point-in-time volumes for which the group-set bookmark failed;

rebuilding the data on the volumes for which the group-set bookmark failed; and

replicating the volumes for which the group-set bookmark failed,

wherein the first quorum is a minimum number of servers available to run a distributed application from a total number of servers, and

wherein a subset of the total number of servers is not available.

2. The method of claim **1**, further comprising:

receiving a request to access a group-set bookmark;

determining if a subset of the available volumes for which the bookmark exists forms a quorum;

allowing access to the groups-set bookmark if the quorum exists; and

denying access to the group-set bookmark if the quorum does not exist.

3. The method of claim **2**, further comprising:

rolling back working consistency groups to an image associated with the group -set bookmark;

determining if the rolling back is successful;

exposing volumes on the target site if the rolling back is successful;

starting an application on servers on the target site if the rolling back is successful; and

undoing rolling back operations if the rolling back is unsuccessful.

4. The method of claim **1**, further comprising:

determining K servers, K servers designating a subset of the servers on the target site with an unavailable journal;

determining M servers, M servers designating a subset of servers on the target site with an accessible journal but the accessible journal does not include the group-set bookmark;

determining if a set of servers denoted by N less the K and M servers form a second quorum, N being the set of servers on the target site;

displaying the group-set bookmark if the second quorum is available; and

displaying K and M servers if the second quorum is unavailable.

5. The method of claim **4**, further comprising indicating to a user if the sets of K or M servers are non-empty.

6. An article comprising:

a non-transitory machine-readable medium that stores executable instructions, the instructions causing a machine to:

determine that a first quorum of servers is available at a production site and a target site;

generate a group-set bookmark if the first quorum of servers is available;

fail the operation of generating the group-set bookmark if there is not a first quorum of servers available;

13

determine L servers, L servers indicating a union of a set of servers on the target site with an unavailable journal or a journal that does not include the group-set bookmark and a union of servers at a production site which are not available; 5

determine if a set of N servers less the L servers forms a quorum, N being the set of servers on the target site; perform failover to currently used servers; allow access to available volumes at the target site which are not part of the quorum; 10

notify an application that volumes are to be rebuilt; start the replication of the volumes which are available also at the production site, wherein starting the replication of the volumes which are available also at the production site comprises replicating from working servers at the target site to the production site; 15

expose at a different point-in-time volumes for which the group-set bookmark failed; rebuild the data on the volumes for which the group-set bookmark failed; and 20

replicate the volumes for which the group-set bookmark failed, wherein the first quorum is a minimum number of servers available to run a distributed application from a total number of servers, and 25

wherein a subset of the total number of servers is not available.

7. The article of claim 6, further comprising instructions causing a machine to:

receive a request to access a group-set bookmark; 30

determine if a subset of the available volumes for which the bookmark exists forms a quorum; allow access to the group-set bookmark if the quorum exists; deny access to the group-set bookmark if the quorum does not exist; and 35

roll back working consistency groups to an image associated with the group-set bookmark; determine if the rolling back is successful; expose volumes on the target site if the rolling back is 40

successful; start an application on servers on the target site if the rolling back is successful; and undo rolling back operations if the rolling back is unsuccessful. 45

8. The article of claim 6, further comprising instructions causing a machine to:

determine K servers, K servers designating a subset of the servers on the target site with an unavailable journal; determine M servers, M servers designating a subset of 50

servers on the target site with an accessible journal but the accessible journal does not include the group-set bookmark; determine if a set of servers denoted by N less the K and M servers form a second quorum, N being the set of servers 55

on the target site; display the group-set bookmark if the second quorum is available; and display K and M servers if the second quorum is unavailable. 60

9. The article of claim 8, further comprising instructions causing a machine to indicate to a user if the sets of K or M servers are non-empty.

10. An apparatus, comprising:

circuitry configured to: 65

determine that a first quorum of servers is available at a production site and a target site;

14

generate a group-set bookmark if the first quorum of servers is available; and fail the operation of generating the group-set bookmark if there is not a first quorum of servers available; determine L servers, L servers indicating a union of a set of servers on the target site with an unavailable journal or a journal that does not include the group-set bookmark and a union of servers at a production site which are not available; 5

determine if a set of N servers less the L servers forms a quorum, N being the set of servers on the target site; perform failover to currently used servers; allow access to available volumes at the target site which are not part of the quorum; 10

notify an application that volumes are to be rebuilt; start the replication of the volumes which are available also at the production site, wherein starting the replication of the volumes which are available also at the production site comprises replicating from working servers at the target site to the production site; 15

expose at a different point-in-time volumes for which the group-set bookmark failed; rebuild the data on the volumes for which the group-set bookmark failed; and 20

replicate the volumes for which the group-set bookmark failed, wherein the circuitry comprises at least one of a processor, a memory, programmable logic and logic gates, wherein the first quorum is a minimum number of servers available to run a distributed application from a total number of servers, and 25

wherein a subset of the total number of servers is not available.

11. The apparatus of claim 10, further comprising circuitry configured to:

receive a request to access a group-set bookmark; determine if a subset of the available volumes for which the bookmark exists forms a quorum; 30

allow access to the group-set bookmark if the quorum exists; deny access to the group-set bookmark if the quorum does not exist; and 35

roll back working consistency groups to an image associated with the group-set bookmark; determine if the rolling back is successful; expose volumes on the target site if the rolling back is 40

successful; start an application on servers on the target site if the rolling back is successful; and undo rolling back operations if the rolling back is unsuccessful.

12. The apparatus of claim 10, further comprising circuitry configured to:

determine K servers, K servers designating a subset of the servers on the target site with an unavailable journal; determine M servers, M servers designating a subset of 50

servers on the target site with an accessible journal but the accessible journal does not include the group-set bookmark; determine if a set of servers denoted by N less the K and M servers form a second quorum, N being the set of servers 55

on the target site; display the group-set bookmark if the second quorum is available; and display K and M servers if the second quorum is unavailable. 60

15

13. The apparatus of claim **10**, further circuitry configured to indicate to a user if the sets of K or M servers are non-empty.

* * * * *

16